



Abstract of the Disclosure

- 5 A middle-tier Web server with a queryable cache that contains items from one or more data sources. Items are included in the cache on the basis of the probability of future hits on the items. The probability determination may be made by an element located in the Web server, at the data source, or elsewhere in the network to which the Web server and the data source belong. The determination may be made by analyzing the hit rates on
- 10 item or by means of information such as a period of time during which frequent hits can be expected. When the data source determines that an item that has been included in the cache has changed, it sends an update message to the server, which updates the item if it is still included in the cache. In a preferred embodiment, the data source is a database system and triggers in the database system are used to generate update messages. In a
- 15 preferred embodiment, the data access layer determines whether a data item required by an application program is in the cache. If it is, the data access layer obtains them from the cache; otherwise, it obtains them from the data source. In the preferred embodiment, the Web application programs use global data set identifiers to identify data sets, while the cached data sets are identified by local data set identifiers. The cache receives a
- 20 query context of global data set identifiers, determines whether the data sets of the query context are cached, and if they are, provides the local data set identifiers for the query context to the access layer, which then uses the local data set identifiers to query the cache.